4-7-08 09/881,505

Please amend the Specification as follows:

Please amend the present title as follows:

Two-~~stage-~~phase commit with queryable caches

Please amend page 6, lines 6-8 as follows:

a1 The two-~~stage-~~phase commit process used for transactions between distributed databases needs to be optimized for the situation where the transaction does not result in modifications to particular ones of the distributed databases.

Please amend page 6, lines 15-24 as follows:

a2 The problem of optimizing the two-~~stage-~~phase commit protocol is solved as follows: in general, when an action is carried out in a distributed system, a component of the distributed system that is involved in the action is a coordinator for the action and other components involved in the action are cohorts for the action. During the action, the cohorts send messages which are available to the coordinator. In the optimization, each cohort augments messages which are available to the coordinator with information which indicates relevant state of the cohort with regard to the action. The coordinator reads the messages and retains the most recent relevant state for each cohort and performs an action according to the relevant state.

AW 4-7-08

Please amend page 6, line30-page 7 line ~~5~~ 2 as follows:

a3 When the protocol is a two-~~stage-~~phase commit protocol, the action performed by the coordinator is sending an *abort* message according to the two-stage commit protocol to a cohort when the retained state for the cohort indicates that the transaction will not modify data in the cohort.

Please amend page 16, lines 5-21 as follows:

a4 In the database context, the two-phase commit protocol works like this: one of the distributed databases is the *coordinator* for the database transaction for which the protocol is being used; the other distributed databases are the *cohorts*. It should be noted here that coordinator and cohort need not be fixed roles. In many cases, the database at which a transaction is initiated